

# The evolutionary design of a Knowledge Network to support knowledge management and sharing for lifelong learning

Patrick McAndrew, Doug Clow, Josie Taylor and James Aczel

*The authors work at The Open University's Institute of Educational Technology. Patrick McAndrew is Head of Centre for Information Technology in Education (CITE). Josie Taylor is Senior Lecturer, and Director of OU's UserLab. Doug Clow is Lecturer in Interactive Media Development, and Deputy Head of CITE. James Aczel is Lecturer in New Technology in Teaching, and Director of the OU Knowledge Network. Address for correspondence: Dr Patrick McAndrew, Institute of Educational Technology, The Open University, Walton Hall, Milton Keynes, MK7 6AA, UK. Tel: +44 (0)1908 652638; email: p.mcandrew@open.ac.uk*

## **Abstract**

Knowledge management (KM) and knowledge sharing are important factors that support lifelong learning, and enable people to continue developing throughout their careers. The concept of a Community of Practice (Wenger, 2000) is attractive in drawing together people whose work shares similar aspects, and consideration is given here to how technology can be used to develop and support such a community.

In this paper, concepts from the Community of Practice literature are used to consider the development of a software environment for people working as a community in the area of lifelong learning. The intention was to design the system in an evolutionary way, using a minimal set of essential elements which would be elaborated according to user feedback. Three key design questions are considered: Who can contribute resources to such a system? What happens to existing practices? How is the community engaged?

We conclude that, in lifelong learning, knowledge management supported by a software environment offers a good way to bring together communities, resources and experience, but to achieve these benefits, great care needs to be exerted in introducing the system, and maintaining existing work practices.

## **Introduction**

### **Background**

The Open University (UK) has been producing undergraduate distance education courses for 30 years. When a new course is to be constructed, a multi-skilled course team is assembled consisting of academics, editors, web designers, producers of audio-visual material, multimedia designers, evaluators and specialists in course design. A vast amount of experience has been accumulated over the years, and new developments in media and teaching continue. In this context, the role of the Institute of Educational Technology (IET) is to ensure that University staff engaged in professional development are able to stay ahead in their field. It does this in a variety of ways, and regards staff in the organisation as lifelong learners, largely because the

body of knowledge is never static, and because professional development is an on-going activity.

Given the quantity of information that the University has collected in the form of reports and papers, computing systems are an obvious support tool to help staff work their way through the resources. IET is well aware, though, that staff do not want only to obtain and read formal documents - they want to know what colleagues are currently thinking, what methods and approaches are currently being used; and they want the opportunity to discuss ideas with colleagues across the University. But no-one has time to attend workshops or other face-to-face events to facilitate these needs.

In response to this situation, in 1999, members of IET formed a team to develop a Knowledge Network (KN). The title was intended to be ambiguous – the software system itself was a network, as were the groups of people it was intended to support. The goal was to examine how individual learners could be supported in forming their own communities of practice through the use of software tools, but in situations where such communities had not yet gelled – the software itself was intended to catalyse the formation of the community. This means that requirements gathering was compromised – the KN team did not know what the community would want, and neither did potential members. In any case, as a community forms and matures, goals might well change. The KN team realised that if the system were to be accepted, they would have to stay as responsive and flexible as possible in order to convince users that their needs were paramount, rather than the needs of either the system itself, or the management concerns of the University.

### ***Communities of Practice***

Wenger, McDermott & Snyder (2002) encourage consideration of various tools that can be associated with a community of practice, as follows:

- A home page to assert their existence and describe their domain and activities
- A conversation space for online discussions
- A repository for their documents, including research reports, best practices and standards
- A good search engine to find things in their knowledge base
- A directory of membership with some information about members' areas of expertise in the domain
- In some cases, a shared workspace for synchronous electronic collaboration, or to enhance teleconferences with visuals
- Community management tools, mostly for the co-ordinator, but sometimes also for the community at large (e.g. knowledge of who is participating actively, which documents are downloaded, which documents need updating, traffic flow, etc.).

However, it is important to recognise that “Knowledge Management” is not a stable, standardised set of deliberate, context-independent processes, each of which requires a support tool. Rather, a particular community's changing knowledge needs should be met unobtrusively, without demanding atypical effort to learn how to use a tool or perform a task on-line. But even when knowledge management is well supported by standardised processes and systems, it will depend for its success on community-

specific goals, resources and practices (Wasko and Faraj, 2000), which indicates a need for any system to stay flexible and responsive to users.

One important principle that we see emerging from this literature and from the software engineering literature is the value of cyclical development and testing with the community the tools are intended to serve. Thus, we suggest, the software development process should not be seen as a distinct practice that is separated from the community, but as a key part of the iterative knowledge sharing activity it aims to support. We therefore argue for evolutionary design, a variation on the rapid prototype and test method. The main difference is that initial versions of the Knowledge Network were not mere prototypes – they were fully functioning systems which were used by groups of workers to achieve their goals.

In addition to the Communities of Practice literature, the design process was also informed by Activity Theory (Engeström, 1987) and a psychological perspective (Aczel, 1998) drawing from the classic approach of Karl Popper. The initial Activity Theory framework has been set out elsewhere (McAndrew and Taylor, 2000). These more theoretical concerns are beyond the scope of this paper and may form the basis of future papers.

As the KN team considered this flexible approach to design, and tried to identify what the needs of the community might be, it became clear that three key questions required attention both from the point of view of potential users, and from the point of view of the institution:

1. Who can contribute to the knowledge in the system?
2. What happens to existing practices?
3. How is the community engaged to participate in knowledge sharing?

## ***The OU Knowledge Network***

The OU Knowledge Network, then, was initially conceived of as a set of tools along the lines of those suggested by Wenger, McDermott & Snyder (op.cit) developed to meet the needs of overlapping sets of people engaged in lifelong learning, as both deliverers of courses and learners themselves. The design was to be informed by explicit decisions on the three questions set out above:

### **1. Who can contribute?**

The list of tools suggested by the Community of Practice literature includes many that potentially allow members of the community to contribute. It would seem obvious, and it seems to be a tacit assumption in the literature, that providing the widest possible opportunity for discussion, iteration and feedback would be best, but practical constraints can mitigate against this – discussions can lose focus, issues become attenuated with too many concerns, and so on. Furthermore, from an institutional perspective, there are potential risks in allowing anyone to contribute. Who would operate a quality assurance process for knowledge being circulated?

Nevertheless, the KN development team took the decision that the system should provide a straightforward means by which any member of the organisation could publish materials and respond to material published by others, with no editorial control or moderation of contribution. This would also avoid the potential bottleneck

that an approval process would necessarily impose, thereby enabling a more rapid build-up of resource.

This decision – “anyone can publish” – was in agreement with the principle that KM tools only achieve high levels of acceptance, trust and productive usage if they fit easily into everyday working practices. It was also thought that the professional members of the community would have no difficulty in identifying weak material, and that it would simply drop out of circulation in due course.

## **2. What happens to existing practices?**

Fundamental to the Community of Practice literature is a respect for the existing knowledge-sharing practices of a community. Obviously, if a set of tools is being introduced there must be some desire to enhance or change existing practices. The second key decision is the degree to which existing practices are supported or supplanted. Rather than seeking to supplant existing practices immediately, the development team decided to support existing practices in the organisation. This required a system that could search multiple websites and data sources simultaneously, to enable those who wished to share their data to do so according to the practices to which they were already accustomed. There were, therefore, no constraints on data sources in the system – information could be found from many sources.

These two decisions – “anyone can publish” and “let a hundred data sources bloom” – constitute a distributed publishing model.

## **3. How is the community engaged?**

No matter how good a set of tools are, though, there is no guarantee that they will be spontaneously taken up and used by the community. The development team decided to follow the evolutionary design-and-test development process set out above in order to engage the community in the Knowledge Network, and its development. In parallel, time was set aside prior to the creation of the system for engaging particular communities in thinking about how they wanted to disseminate their work using electronic methods.

This “hearts-and-minds” work was aimed at ensuring that by the time the first full version of the system arrived, not only would the dissemination mechanisms have been shaped by the deliberation, but that the knowledge cultures of the various communities would have changed from one of hoarding (because “knowledge is power”) to one of sharing (because shared knowledge leads to better, more useful knowledge). This shift in knowledge culture was necessary if the strategy of distributed publishing was to work.

### **Supporting and enhancing existing practices**

Having worked alongside these lifelong learners, the KN team had identified that the clear focus of their activities was access to existing and new materials, a task in which they were often frustrated. The strategy to tackle this problem is to collect the materials in an easy-to-search data store. The tool used was a browser-based front-end to a database organised so that the database appears to be a single dataset to end-users, though the search engine is actually seamlessly searching a range of data sources behind the scenes. Furthermore, materials are automatically cross-referenced, so that users can see which other materials relate to the item they are viewing.

The “anyone can publish” decision enabled a rapid transfer of existing documents to the system, and coupled with the decision on supporting existing practices and the “hearts and minds” work resulted in a rapid take-up of this facility. A formative evaluation (Twining & Rico, 2002) found that the majority of users using the Knowledge Network were using it to locate documents. Moreover, this database function of the Knowledge Network appeared to constitute a large part of its perceived value: the study found that the Knowledge Network saved users’ time looking for materials, helped them locate people with the knowledge they required, and helped them find information.

Within a year, around 20% of OU staff were using the Knowledge Network; and within two years, around 40% of OU staff were using it. Interview evidence suggested that locating information was the main purpose of users as a whole, and that they were mainly successful. One of the respondents to the study said:

“The Knowledge Network is a useful way to find out what people in the OU have already found out about teaching issues, especially 'new' issues concerning e.g. use of technology. A good place to find contacts, published reports and avoid 'reinventing the wheel'.”

The experience of this case study would suggest that, despite it being much less glamorous than other e-learning tools, a shared database combining an uncomplicated, familiar interface with a powerful search engine and rich content is arguably one of the most useful knowledge management tools one can provide for motivated independent learners.

## Sharing Materials

How effective was the distributed publishing strategy in meeting the knowledge *sharing* needs of the particular lifelong learners under consideration?

Participants in the evaluation a year after launch (Twining & Rico, op cit) consistently reported that the Knowledge Network helped them disseminate their work and to manage their own documents. “Early adopters” of the system might tend to be more tolerant, and more likely to be involved in publishing than other users, so an initial large overlap between publishers and accessors was expected. Data from two years after launch suggests that this overlap had decreased.

Users who share materials interact with the subject matter in different ways. As was the case for access to materials via the database, it is clear from a subsequent study that the various communities that publish materials are not homogeneous. The data show that the area initially targeted for engagement with the system was responsible for the majority of the output; other distinct groups provided smaller collections of data or individual reports. The system worked well to bring together different report series in one place, and to allow users to see the smaller, more diverse set of reports from other departments and from fellow learners alongside the major report series.

The development of the Knowledge Network demonstrates that a distributed publishing strategy can enable lifelong learners to access knowledge that was previously hard to access. The evidence of benefits when learners share their own materials is less clear. More research is needed on this important aspect of KM for lifelong learners. A priori, one would expect that simply working towards a goal of putting work into a form that can be shared with colleagues would have value, and more so if discussions of that work follow.

## **Collaboration Tools**

In addition to the database, further tools were provided to learners, and we now consider how the various communities exploited these tools.

The Knowledge Network provides technology for users to create web sites without technical knowledge. The web sites have relatively sophisticated features such as access control, discussion, bulletin boards, news, forms, automatic cross-referencing, search, hit count statistics, and subscriptions (receiving an email notification when the selected resource is updated). These community websites were termed “KN workspaces”. So how effective were these collaborative workspaces in meeting the knowledge needs of these lifelong learners?

Those participants in the evaluation study a year after launch who used the workspaces reported that their collaboration was supported. Overall, though, very few users used them and fewer than 10% of all users contributed to discussions. An online questionnaire found that the vast majority of staff either did not know about or did not understand how to use collaborative workspaces. A training programme was started to help people understand the collaboration facilities and how they could be used to support their work.

In addition, a series of structured activities were created, to help engage the target population of lifelong learners with issues in particular areas of interest (teaching and learning online, and creating quality courses). These blended face-to-face, self study, discussion and resource descriptions within a clear set of time-limited activities.

Meanwhile, several web sites that had been previously independent of the Knowledge Network were moved to be hosted by these workspaces. This not only provided the benefits of more scaleable content management and of the collaboration facilities, but also introduced key users to the potential of the technology. Some of these websites are public. For example, KN technology powers a national library to support good practice in the re-use of educational software, and a higher education network, with members in 160 institutions from 18 countries (the Reusable Educational Software Library, RESL - <http://www.resl.ac.uk>).

This effort increased the usage of the collaborative tools dramatically. A year after the first formal evaluation (i.e. two years after launch), there were over 300 workspaces. With over 1500 web pages created, and over 25,000 visits in total, and about 2000 page impressions every week (with a user-base of about 7000), the workspaces do appear to be of value to the target learners.

The evidence is that the communities are diverse, but that workspaces fall into one of the following types, with a few combining elements of more than one type:

- A shared private work area
- A dissemination website
- An authoritative overview of a topic
- A learning activity

What is common to all these types of workspace is that the extent of discussion and engagement in any given workspace is extremely variable, and that drivers to create the workspace tend to come from the users’ own needs and motivations, rather than from the technical system itself.

## **Conclusions**

This case study has shown that the original design decisions, based upon the simple 3-issue model, were sufficient to support communities of self-directed learners within an organisational context. The tools were found to be useful, and uptake was surprisingly high given that the only incentive was the tool itself – there was no requirement for anyone to use the system. Interestingly, the response of the larger organisation was mixed. The local context of the Institute was supportive of the development of the Knowledge Network, as its staff had long appreciated the need for effective communication and sharing with busy colleagues. However, in the wider context of the University, the Knowledge Network was viewed as an anarchic threat to existing new systems of document management. The very features that were most prized in the KN (the ability for anyone to publish, the absence of editorial control, the freedom to exchange with many different kinds of learners) were seen as its biggest flaws. A serious amount of the KN team's time was spent in promoting the difference between the KN and formal methods of document storage – the KN team wanted to preserve the dynamic flexibility inherent in sharing, rather than construct the definitive body of knowledge that could be codified.

In this paper we have considered the growing use of knowledge management in support of the knowledge sharing process, and the use of a software system to support learning communities of practice in this – the Knowledge Network. In doing so, we have understood more about the power of the concept – i.e. that knowledge sharing is a key component of the formation of operation of lifelong learners as a community – and we have illustrated particular methods to achieve the sharing. The characteristics that have enabled the success we have achieved with the KN are centred on ease of use and integration with an environment: the tools need to lower the barrier towards sharing rather than become an end in themselves. Of greater importance than the tools, though, is the link to patterns of working and the care with which the concepts are introduced.

Our three key design questions were: Who can contribute resources to such a system? What happens to existing practices? How is the community engaged? These proved sufficient to drive a development process, both in terms of software development, and in terms of the activities of our target users, that has resulted in a rich working environment to share knowledge and experience. Resisting organisational pressure to formalise the system and its processes also has demonstrated to the user community that the developers were anxious to respond to their actual needs, rather than impose methods on them. This has led to a sense of closer community amongst the groups using the KN, and underlines the value of keeping tightly focused on the needs of the users.

## **References**

- Aczel J C (1998) *Learning Equations using a Computerised Balance Model: A Popperian Approach to Learning Symbolic Algebra* DPhil Thesis, University of Oxford.
- Engeström Y (1987) *Learning by Expanding: An Activity-Theoretical Approach to Developmental Research* Orienta-Konsultit Oy, Helsinki, Finland.

McAndrew P and Taylor J (2000) Organising for a Knowledge Network, *HCI'2000*, Conference of the British Computer Society, 5 – 8 September 2000, Sunderland University, UK.

Twining P and Rico M (2002) *Knowledge Network Evaluation – Final Report*. Retrieved online 28/5/2004 at <http://kn.open.ac.uk/public/document.cfm?documentid=2627>

Wasko M M and Faraj S (2000) 'It is what one does': why people participate and help others in electronic communities of practice *Journal of Strategic Information Systems* 9 2-3, 155-173.

Wenger E (2000) *Communities of Practice: Learning, Meaning, and Identity* Cambridge University Press, Cambridge, UK.

Wenger E C, McDermott R and Snyder W C (2002) *Cultivating Communities of Practice: A Guide to Managing Knowledge* Harvard Business School Press, Cambridge, MA.